

**METHOD FOR TRACKING TRANSMISSION STATUS OF DATA TO
ENTITIES SUCH AS PEERS IN A NETWORK**

TECHNICAL FIELD

This application relates, in general, to routers, and more particularly to methods for tracking a transmission status of data to be communicated to entities such as peers in a network.

BACKGROUND

Typically, routers have a route processor which communicates with other routers in a network to obtain updates of route information regarding the network. A route processor maintains a routing table which contains a large number of entries or elements for use by the router to determine how to forward packets that are received by the router. Conventional routing tables may contain hundreds of thousands of routing entries or elements, and as a greater number of devices such as personal data assistants (PDAs), cell phones, and other devices are assigned IP addresses, the number of entries in routing tables is expected to grow substantially over the coming years.

In the Internet, a routing protocol known as Border Gateway Protocol (BGP) provides a way for routers to share routing information with other routers throughout the network. These other routers that a router speaks to are called "BGP peers." Fig. 1 illustrates an example of a BGP speaking router 20 in communications with three BGP peers (routers) 22, 24, 26 shown as Peer 1, Peer 2 and Peer 3 in this example. As the router 20 learns new routing information, it may need to communicate this new information to one or more of its peers 22, 24, 26.

Fig. 2 shows one method to track a transmission status of data elements 30, such as routing entries (shown as I1 to I4), with respect to peers 32 (i.e., which peers need to be updated with particular data elements). In Fig. 2, data elements I1-I4 are shown as column headers, and Peers 1-3 are shown as row headers. Such a table data structure 34 can be used in a router to track which peers need to be updated with

1 what data elements. For example in Fig. 2, data elements I1 and I4 need to be
2 communicated to Peer 1, data elements I2 and I4 need to be communicated to Peer 2,
3 and data elements I1 and I4 need to be communicated to Peer 3, in this example.

4 As recognized by the present inventors, such a table data structure 34 to track
5 the status of data elements 30 as shown in Fig. 2 is undesirable because it requires
6 large amounts of memory to form this data structure. This is particularly true where
7 the number of data elements 30 may be on the order of hundreds of thousands of
8 routing entries or data elements, such as in the case of route entries in a routing table
9 of a router which are to be communicated to numerous peers in the network.

10 As recognized by the present inventors, what is needed is a method for
11 tracking the state or transmission status of data elements with respect to peers or
12 other devices or other entities in a network.

13 It is against this background that various embodiments of the present
14 invention were developed.

15 SUMMARY OF THE INVENTION

16 In light of the above and according to one broad aspect of one embodiment of
17 the present invention, disclosed herein is a method for tracking a transmission status
18 of one or more data elements or routing entries to one or more devices. In one
19 example, the method includes providing a list including one or more devices and one
20 or more data elements; processing the list to determine a data element of the one or
21 more data elements to transmit to a device of one of the one or more devices; and
22 upon successfully transmitting the data element to the device, adjusting the list so
23 that the list indicates that the device has received the transmitted data element. In
24 this manner, the status of whether a particular data element has been sent to a
25 particular device can be derived from the list.

26 The methods described herein are particularly useful in the context of network
27 routers, where the number of data elements in a routing table of a router may be
28 substantial. In one embodiment, the list may be implemented as a linked list,
29 wherein data elements or devices may be added or deleted from the list. The list may

1 also include a global version number, a local version number associated with each
2 data element in the list, and a local version number associated with each device in
3 the list. The local version number associated with a data element in the list may be
4 set to a value of the global version number at a time when the data element was
5 added to the list. The local version number associated with a device in the list may
6 be set to an initial value of zero and may be reset to the local version number of a
7 data element after the data element is successfully transmitted to the device.

8 Processing the list may include the operations of locating a device in the list
9 which is nearest to a start of the list; obtaining the version number for the device;
10 and comparing the version number to the global version number to determine if the
11 device should have a data element transmitted to the device. If the version number
12 of the device is not equal to the global version number, then in one example the
13 device should have a data element transmitted to the device.

14 In one embodiment, the operation of adjusting the list includes repositioning
15 the device within the list adjacent to the data element and closer to an end of the list
16 than the data element and resetting the local version number of the device to be equal
17 to the local version number of the data element transmitted to the device.

18 According to another broad aspect of another embodiment of the invention,
19 disclosed herein is a method for transmitting one or more data elements to one or
20 more devices. In one example, the method includes providing a list including one or
21 more devices and one or more data elements; processing the list to determine a
22 data element of the one or more data elements to transmit to a device of one of the
23 one or more devices; transmitting the data element to the device; and adjusting the
24 list to indicate that the device has received the transmitted data element.

25 According to another broad aspect of another embodiment of the present
26 invention, disclosed herein is a router. In one example, the router include a module
27 for providing a list including one or more devices and one or more data elements; a
28 module for processing the list to determine a data element of the one or more data
29 elements to transmit to a device of one of the one or more devices; a module for

1 transmitting the data element to the device; and a module for adjusting the list so
2 that the list indicates that the device has received the transmitted data element.

3 Various embodiments of the present invention may be embodied as computer
4 program products including computer usable medium and computer readable code
5 embodied on the computer usable medium, the computer readable code including
6 computer readable program code devices configured to cause the computer to
7 perform or effect one or more of the operations described herein.

8 The features, utilities and advantages of various embodiments of the invention
9 will be apparent from the following more particular description of embodiments of
10 the invention as illustrated in the accompanying drawings.

11 **DESCRIPTION OF THE DRAWINGS**

12 Fig. 1 illustrates a block diagram of a router communicating with a plurality
13 of peers.

14 Fig. 2 illustrates an example a table data structure for tracking a transmission
15 status of data elements with respect to peers or other devices or other entities in a
16 network.

17 Fig. 3 illustrates an example of logical operations for forming a transmission
18 list for tracking the transmission status of data elements with respect to peers or
19 other devices or other entities in a network, in accordance with one embodiment of
20 the present invention.

21 Fig. 4 illustrates an example of a transmission list, in accordance with one
22 embodiment of the present invention.

23 Fig. 5 illustrates an example of logical operations for transmitting elements to
24 the members in a transmission list, in accordance with one example of the present
25 invention.

26 Fig. 6 illustrates an example of logical operations for adding a new member,
27 in accordance with one embodiment of the present invention.

28 Fig. 7 illustrates an example of logical operations for adding a new element to
29 a transmission list, in accordance with one embodiment of the present invention.

Fig. 8 illustrates an example of logical operations for deleting a member, in accordance with one embodiment of the present invention.

Fig. 9 illustrates an example of logical operations for deleting an element from a transmission list, in accordance with one embodiment of the present invention.

Fig. 10 illustrates an example of logical operations for modifying an element in a transmission list, in accordance with one embodiment of the present invention.

Figs. 11A-11N illustrate examples of transmission lists changing over time as elements are added to the transmission lists, and as elements are transmitted to members, in accordance with one embodiment of the present invention.

Figs. 12A-12B illustrate an example of a use of a marker with a transmission list, in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION

According to one broad aspect of one embodiment of the present invention, disclosed herein is a method and a data structure for tracking the transmission status of a plurality of data elements or routing entries to a plurality of intended receivers, devices, modules, peers or entities, referred to as members. The methods and data structures described herein are particularly useful in the context of network routers such as shown in Fig. 1, where the number of data elements or routing entries in a routing table of a router may be substantial. In the router of Fig. 1, a transmission list, described below, provides information to a BGP application program running in the router so that the router updates its peer members with new route information or other data as needed. In one example, the transmission list of an embodiment of the present invention is used to track the state of which peer members have received certain data elements. One or more steps of the methods disclosed herein can be used to transfer large routing databases between routers in a network. Various embodiments of the present invention will now be described.

1 Referring to Fig. 3, Fig. 3 illustrates an example of logical operations for
2 tracking a transmission status of the plurality of data elements (i.e., routing entries)
3 with regard to a plurality of members (i.e., intended receivers, devices, modules,
4 peers, entities, etc.). At operation 40, a list is established with the members and data
5 elements. In one embodiment, the list is implemented as a linked list wherein a
6 plurality of data elements and members are linked together within the list, and the
7 relative positions of the members and elements within the list are used to derive
8 transmission status or state information. An example of such a linked list 42 is
9 illustrated in Fig. 4, wherein data elements 44, shown as I1, I2, I3, I4 are in a linked
10 list with members 46, shown as P1, P2, and P3, in this example. It is understood that
11 while Fig. 4 illustrates only three members 46 and four data elements 44, the number
12 of members 46 and data elements 44 in the list 42 could be substantially larger in
13 order to accommodate, for instance, a large number of route entries in a routing table
14 of a router.

15 At operation 48 of Fig. 3, the list is processed to determine what elements
16 need to be sent to which members. For instance, a router can utilize the list to
17 examine and quickly determine which particular data elements need to be sent to
18 particular members. At operation 50, the list is adjusted to reflect when a particular
19 member has been sent particular elements. In other words, when a member has been
20 sent a data element, the transmission list is adjusted to reflect this fact. In one
21 embodiment, operations 48-50 are implemented in a continuous loop which may be a
22 background process running in a router. Various examples of transmission lists in
23 accordance with the present invention will be described in greater detail below.

24 In one embodiment, a transmission list is maintained and may be implemented
25 in the programming language of C as a data structure or linked list. As shown in Fig.
26 4, a transmission list 42 may include one or more members 46 – which represent
27 entities or routers which are to receive data – and one or more elements 44 – which
28 represent data elements, data values, routing entries or information to be transmitted
29 to the one or more members 46. In one example, the members 46 and elements 44

1 are arranged in a linked list 42, wherein each item in the list references the item
2 adjacent to it in the list. In one example, the links between elements or members in
3 the transmission list are bi-directional.

4 The transmission list 42 can be used to maintain the state of transmission
5 status of the one or more elements 44 to the one or more members 46. In one
6 embodiment, if an element 44 is positioned closer to the end of the transmission list
7 42 than a member 46, such relative positioning indicates that the member 46 has not
8 yet received the element 44. Hence, at an initial time prior to any transmissions, the
9 members 46 would be located at the start of the list 42, and the data elements 44
10 would be positioned in the list 42 approximate the end of the list, as shown in Fig. 4.

11 Conversely, if a member 46 is positioned closer to the end of the transmission
12 list 42 than a data element 44, then such positioning indicates that the member 46
13 has already received the element 44. Hence, at steady state when all members 46
14 have received all data elements 44, then all of the members 46 would be positioned
15 within the list approximate the end of the transmission list, and all of the elements
16 would be positioned within the list approximate the start of the transmission list.

17 As shown in Fig. 4, a pointer 52 may be maintained indicating the start or
18 beginning of the transmission list, and another pointer 54 may be maintained
19 indicating the end of the transmission list. Further, a global version number 56 may
20 be maintained based on the elements in the list, and as elements are added to the list,
21 the global version number 56 may be incremented. Each member 46 is provided
22 with its own local version number 58 (abbreviated as "v" in Fig. 4) which, in one
23 example, is initialized to zero at the time when the member 46 is added to the
24 transmission list. Further, in one example, each element 44 may be provided with a
25 local version number 60 (abbreviated as "v" in Fig. 4) which is initially set to a
26 value of the global version number 56 when the element was added to the
27 transmission list. In one example, when an element needs to be added to the list, the
28 global version number is incremented and the element's version number is set to this

1 incremented global version number. The element is then added to the end of the
2 transmission list, in one example.

3 By utilizing a transmission list, flow control is provided in that the data
4 elements 44 may be transmitted to members 46 in a graceful and ordered fashion.
5 Embodiments of the present invention may be particularly beneficial in networking
6 wherein the transmission list 42 may include hundreds of members 46 – such as
7 “peer” routers – and there may be hundreds of thousands of elements 44 – such as
8 routing information including route updates. By the use of a transmission list, the
9 computing operations to determine the next data elements to be transmitted to the
10 next member can be performed in constant time, as opposed to a matrix system of
11 Fig. 2 which may be a time intensive search operation to locate such information.
12 Further, in contrast to a system utilizing a matrix or table such as shown in Fig. 2,
13 wherein the amount of transient memory utilized by such a data structure is
14 approximately proportional to the number of members multiplied by the number of
15 elements, the amount of memory utilized by a transmission list 42 in accordance with
16 an embodiment of the invention is proportional to the number of members plus the
17 number of elements multiplied by the amount of memory to maintain the linkage in
18 the list, such as four bytes. Accordingly, as members 46 or elements 44 are added to
19 the transmission list 42, the amount of memory utilized by a transmission list 42
20 would grow linearly, not by a multiplication factor.

21 A member list may also be maintained which includes all of the members 46
22 in the transmission list. By maintaining a member list separate from the
23 transmission list 42, the members in the transmission list can be easily located,
24 particularly where the members are amongst hundreds of thousands of elements in
25 the transmission list.

26 Fig. 5 illustrates an example of logical operations for processing a
27 transmission list to determine a next element to be transmitted to a member, in
28 accordance with one embodiment of the present invention. Referring to Fig. 5, at
29 operation 62, the start of the transmission list is examined and the nearest member is

1 located, in one embodiment. In one example, the locate operation 62 may be
2 performed using a radix search, a patricia tree search, a hash table search, a skip list
3 search, a red-black tree search, or other look-up methodologies for locating an item
4 in a data structure. Alternatively, locating a member may merely involve a walk of
5 the linked list, and for flow control purposes, members could be skipped temporarily
6 and other members selected that are ready for receiving routing information.

7 At operation 64, upon locating a member, the version number associated with
8 the located member is compared with the global version number. Operation 66
9 determines if the member's version number matches the global version number, and
10 if so, this indicates that this member has received all data elements (see block 68)
11 and in one embodiment, control is returned to operation 62. In this instance,
12 operations 62-66 will repeat until a new data element is added or modified such that
13 operation 66 detects that the global version number and a member's version number
14 do not match.

15 Assuming that a member's version number and the global version number do
16 not match, then control is passed to operation 70. Operation 70 locates a next
17 element in the transmission list to transmit to a member. In one embodiment,
18 operation 70 locates the next element nearest to the start of the transmission list. In
19 one example, the element that occurs in the linked list after the member, is the
20 element to be sent to the member. The locate operation 70 may be performed using a
21 non-key based search or using a radix search, a patricia tree search, a hash table
22 search, a skip list search, a red-black tree search, or other look-up methodologies for
23 locating an item in a data structure.

24 At operation 72, the element located by operation 70 is transmitted to a
25 member, preferably the member located by operation 62. The transmission operation
26 72 may be a transmission of a message over a network to another device such as a
27 peer, or may include transmitting or writing data to another process within a device.
28 Operation 74 determines whether the transmission operation 72 was successful and if
29 not, operation 76 retries the transmission operation. In one example, operation 74

1 awaits receipt of an acknowledgment response from the member to which the
2 element was transmitted by operation 72, and when such acknowledgement is
3 received, operation 74 passes control to operation 78.

4 At operation 78, assuming that the transmission operation 72 was successful,
5 operation 78 updates the member's version number to be the same as the version
6 number from the elements which was transmitted to the member. At operation 80,
7 the member is moved within the transmission list to indicate that the member has
8 received the element. In one example, operation 80 positions the member adjacent to
9 the element such that the member is positioned closer to the end of transmission list
10 than the position of the element.

11 Operation 80 may pass control to operation 62 such that these operations may
12 be repeated as needed so that all data elements are transmitted to all members, and as
13 data elements are modified or added, or members are added, the operations of Fig. 5
14 ensure that the elements are transmitted to the appropriate members as needed.

15 The operations of Fig. 5 are beneficial in that they provide a graceful and
16 controlled transmission of the data elements to the members, which is particularly
17 beneficial in networking because some routers can only absorb a finite amount of
18 information at any given time relating to route updates or other routing information.

19 In accordance with embodiments of the present invention, new members can
20 be added to or deleted from a transmission list, and new elements can be added,
21 deleted, or modified. Figs. 6-10 relate to these operations. In Fig. 6, when a new
22 member is added, operation 90 adds the member to the end of the member list and, in
23 one embodiment, the member is added into the transmission list approximate the
24 start of the transmission list before any elements, which, in one example, may be at
25 the start of the transmission list. At operation 92, the new member's version number
26 is initialized, and in one example, is initially set to zero.

27 In Fig. 7, when a new element is added to the transmission list, operation 100
28 adds the new elements to the end of the transmission list. At operation 102, the
29 global version number is incremented to reflect the fact that the new element has

1 been added to the transmission list, and at operation 104, the version number of the
2 element is set to be the value of the global version number.

3 In Fig. 8, when a member is deleted, operation 110 removes the member from
4 the member list and from the transmission list. The effect of operation 110 is that a
5 void or disconnect is created in the member list and the transmission list, and
6 operation 112 reconnects each of these lists such that the voids in the list are patched
7 by referencing to one another the items which were adjacent to the removed member.

8 In Fig. 9, when an element is deleted from the transmission list, at operation
9 120, the element is removed from the transmission list, which creates a void in the
10 list. Operation 122 reconnects the list by referencing to one another the items which
11 were previously adjacent to the removed element. At operation 124, the global
12 version number is adjusted to account for the fact that an element in the list has been
13 deleted. For example, if the element deleted from the list was the highest numbered
14 element in the list wherein the element's version number was equal to the global
15 version number, then operation 124 would decrement the global version number
16 upon deletion of that element. If the deleted element had a version number which
17 was less than the global version number, then upon deletion of that element, the
18 global version number would not need adjustment, in one example.

19 Fig. 10 illustrates an example of logical operations for modifying an element
20 in a transmission list, in accordance with one embodiment of the present invention.
21 A modification of a data element can be implemented as a deletion of the old version
22 of the element and an addition of the new version of the element. At operation 130,
23 the old version of the element is deleted. Operation 130 may be implemented using
24 the operations of Fig. 9, in one example. At operation 132 of Fig. 10, the new
25 version of the element is added to the end of the transmission list (which may be
26 implemented using the operations of Fig. 7). The operations of Fig. 7 take care of
27 incrementing the global version number.

28 In Figs. 11A-N, various examples of a transmission list 140 and a member list
29 142 are illustrated. In Fig. 11A, the transmission list 140 begins with three members

1 P1, P2, and P3, each having local version numbers 144 initialized to zero (shown as
2 v=0). The global version number 146 (shown as "GVN") is also initialized to zero.
3 As indicated in Fig. 11A, a start pointer 148 points to member P1 and end pointer
4 150 points to member P3 of the transmission list. In Fig. 11B, an element I1 is
5 added, and accordingly the global version number 146 is incremented to have a value
6 of one and the element I1 is set to have a version number 152 equal to the global
7 version number 146 of one. The element I1 is added to the transmission list at the
8 end of the transmission list. In Fig. 11C, another element I2 is added to the end of
9 the transmission list, and accordingly, the global version number 146 is incremented
10 to a value of two and the version number 154 associated with the element I2 is also
11 set to have a value equal to the global version number 146 of two. In Fig. 11D,
12 another data element I3 is added to the end of the transmission list, and accordingly
13 the global version number 146 is incremented to a value of three, and the version
14 number 156 of the element I3 is set to the value of the global version number 146 of
15 three.

16 In examining Fig. 11D, it can be seen that, pursuant to example operations of
17 Fig. 5, the transmission list 140 indicates that element I1 should be transmitted to
18 member P1, in one example. In Fig. 11E, assuming that the element I1 is
19 successfully transmitted to member P1, then the member P1 is repositioned within
20 the transmission list 140 adjacent to element I1 and closer to the end of the
21 transmission list. Further, the version number 144 associated with member P1 is set
22 to the value of the version number 156 of the element I1 which was successfully
23 transmitted to member P1 (i.e. to a value of one in this example).

24 In examining Fig. 11E, the transmission list 140 indicates that element I1 is to
25 be transmitted to member P2. In Fig. 11F, the data element I1 is attempted to be
26 transmitted to member P2, however, assuming such attempted transmission was
27 unsuccessful, then the position of member P2 relative to element I1 remains
28 unchanged in the transmission list 140, in one embodiment. Assuming that the
29 transmission of element I1 to member P2 is retried and successful, then as shown in

1 Fig. 11G, member P2 is repositioned in the transmission list 140 adjacent to element
2 I1 closer to the end of the transmission list, and member P2's version number 144 is
3 updated.

4 In examining Fig. 11G, the transmission list 140 indicates that element I1 is to
5 be transmitted to member P3. In Fig. 11H, assuming that the element I1 is
6 successfully transmitted to member P3, then the member P3 is repositioned within
7 the transmission list 140 adjacent to element I1 and closer to the end of the
8 transmission list. Further, the version number 144 associated with member P3 is set
9 to the value of the version number 152 of the element I1 which was successfully
10 transmitted to member P3 (i.e. to a value of one in this example).

11 In examining Fig. 11H, the transmission list 140 indicates that element I2 is to
12 be transmitted to member P3. In Fig. 11I, assuming that the element I2 is
13 successfully transmitted to member P3, then the member P3 is repositioned within
14 the transmission list 140 adjacent to element I2 and closer to the end of the
15 transmission list. Further, the version number 144 associated with member P3 is set
16 to the value of the version number 152 of the element I2 which was successfully
17 transmitted to member P3 (i.e. to a value of two in this example). Note that in Figs.
18 11I to 11N, the member list 142 shown in Figs. 11A-H is not shown which
19 simplifies the drawings.

20 In examining Fig. 11I, the transmission list 140 indicates that element I2 is to
21 be transmitted to member P2. In Fig. 11J, assuming that the element I2 is
22 successfully transmitted to member P2, then the member P2 is repositioned within
23 the transmission list 140 adjacent to element I2 and closer to the end of the
24 transmission list. Further, the version number 144 associated with member P2 is set
25 to the value of the version number 152 of the element I2 which was successfully
26 transmitted to member P2 (i.e. to a value of two in this example).

27 In examining Fig. 11J, the transmission list 142 indicates that element I2 is to
28 be transmitted to member P1. In Fig. 11K, assuming that the element I2 is
29 successfully transmitted to member P1, then the member P1 is repositioned within

1 the transmission list 140 adjacent to element I2 and closer to the end of the
2 transmission list. Further, the version number 144 associated with member P1 is set
3 to the value of the version number 152 of the element I2 which was successfully
4 transmitted to member P1 (i.e. to a value of two in this example).

5 In examining Fig. 11K, the transmission list 140 indicates that element I3 is to
6 be transmitted to member P1. In Fig. 11L, assuming that the element I3 is
7 successfully transmitted to member P1, then the member P1 is repositioned within
8 the transmission list adjacent to element I3 and closer to the end of the transmission
9 list. Further, the version number 144 associated with member P1 is set to the value
10 of the version number 152 of the element I3 which was successfully transmitted to
11 member P1 (i.e. to a value of three in this example).

12 In examining Fig. 11L, the transmission list 140 indicates that element I3 is to
13 be transmitted to member P2. In Fig. 11M, assuming that the element I3 is
14 successfully transmitted to member P2, then the member P2 is repositioned within
15 the transmission list adjacent to element I3 and closer to the end of the transmission
16 list. Further, the version number 144 associated with member P2 is set to the value
17 of the version number of the element I3 which was successfully transmitted to
18 member P2 (i.e. to a value of three in this example).

19 In examining Fig. 11M, the transmission list 140 indicates that element I3 is
20 to be transmitted to member P3. In Fig. 11N, assuming that the element I3 is
21 successfully transmitted to member P3, then the member P3 is repositioned within
22 the transmission list adjacent to element I3 and closer to the end of the transmission
23 list. Further, the version number 144 associated with member P3 is set to the value
24 of the version number 152 of the element I3 which was successfully transmitted to
25 member P3 (i.e. to a value of three in this example).

26 Because in Fig. 11N the version numbers 144 of the members is equal to the
27 global version number 146, this condition indicates that all of the members have
28 received all of the elements. Further, the positions of the members relative to the
29 elements in the list also indicate that all of the members have received all of the

1 elements. Accordingly, it can be seen that by maintaining a transmission list
2 including members and elements, the transmission list tracks and indicates the status
3 as elements are transmitted to members.

4 In another embodiment, one or more markers or indexes may be used with the
5 transmission list to represent or point to the positions of members in the
6 transmission list. In one example, the number of markers is less than or equal to the
7 number of members in the list. In this manner, the markers eliminate the need to
8 "walk down" the transmission list to find a member. For example, where a
9 transmission list is lengthy and contains hundreds of thousands of items including
10 members and elements, use of one or more markers permit operation 62 of Fig. 5 to
11 instantaneously locate the nearest member in a lengthy transmission list. Examples
12 of use of a marker are shown in Figs. 12A and 12B, wherein in Fig. 12A a marker
13 160 is shown representing the location of members P1, P2, and P3 in the
14 transmission list 162. In Fig. 12B, assuming that element I1 was sent to member P1,
15 then a first marker 170 indicates the location of member P1 in the transmission list
16 162, while a second marker 172 indicates the location of members P2, P3 in the
17 transmission list.

18 Assuming that there are 100 members and 100,000 elements, it can be seen
19 that a table structure as shown in Fig. 2 would require an amount of memory
20 proportional to $(1 \text{ bit} \times 100 \times 100,000) = 10,000,000$ bits, whereas a system
21 employing an embodiment of the present invention may use, in one example, an
22 amount of memory proportional to $(32 \text{ bits} \times (100 + 100,000)) = 3,203,200$ bits, plus
23 some small amount of memory for the pointers and/or markers and/or other list
24 management values. Hence, in this example, the amount of memory used by
25 embodiments of the present invention is substantially smaller and can be searched
26 much more quickly than the system of Fig. 2. Furthermore, embodiments of the
27 present invention require relatively less CPU cycles to determine the next element to
28 be transmitted since this can be done in a constant time operation rather than an
29 operation with a matrix model. If an embodiment of the present invention is

1 employed within a router, the amount of memory savings and improvements in
2 computational efficiency will be significant over the system of Fig. 2.

3 The methods disclosed herein can be used to transfer large routing databases
4 between routing peers. While the examples shown and described herein include
5 three members and three elements in a transmission list, it is understood that
6 embodiments of the present invention may include multiple members and multiple
7 elements, such that the transmission list may have hundreds of thousands, or more,
8 of items therein. Also, it is understood that the example operations of Fig. 5 and the
9 other figures herein could be altered to process items in the transmission list in a
10 different manner, while still using the transmission list to indicate and track the
11 status of members and elements in the list. The operations of Fig. 5 are by way of
12 example only, and the variations in the operations – such as reversing the order of
13 one or more of the operations, reversing the orientation of the transmission list,
14 performing a hybrid process, and/or performing a complementary process – are
15 within the scope the invention.

16 One or more operations described herein may be implemented as one or more
17 modules, or portions thereof, in a router, for example as shown in Fig. 1.
18 Furthermore, various embodiments of the present invention may be embodied as
19 computer program products – such as removable media, disks, CD ROMS, or other
20 computer program products – including computer usable medium and computer
21 readable code embodied on the computer usable medium, the computer readable code
22 including computer readable program code devices configured to cause the computer
23 to perform or effect one or more of the operations described herein. Embodiments of
24 the present invention may also be embodied in computer readable code which may be
25 transmitted over a network as packet data, as a carrier wave, or as a propagated
26 signal.

27 While the methods disclosed herein have been described and shown with
28 reference to particular operations performed in a particular order, it will be
29 understood that these operations may be combined, sub-divided, or re-ordered to

1 form equivalent methods without departing from the teachings of the present
2 invention. Accordingly, unless specifically indicated herein, the order and grouping
3 of the operations is not a limitation of the present invention.

4 While the invention has been particularly shown and described with reference
5 to various embodiments thereof, it will be understood by those skilled in the art that
6 various other changes in the form and details may be made without departing from
7 the spirit and scope of the invention.

8